arXiv:quant-ph/0112086v1 15 Dec 2001

Quantum lower bounds for the collision and the element distinctness problems [1]

Yaoyun Shi

Institute for Quantum Information
California Institute of Technology
Pasadena, CA 91125, USA
E-mail: shiyy@cs.caltech.edu

Abstract

Given a function f as an oracle, the collision problem is to find two distinct inputs i and j such that $f(i) = f(j)$, under the promise that such inputs exist. Since the security of many fundamental cryptographic primitives depends on the hardness of finding collisions, quantum lower bounds for the collision problem would provide evidence for the existence of cryptographic primitives that are immune to quantum cryptanalysis.

In this paper, we prove that any quantum algorithm for finding a collision in an r-to-one function must evaluate the function $\Omega\left((n/r)^{1/3}\right)$ times, where n is the size of the domain and $r|n$. This improves the previous best lower bound of $\Omega\left((n/r)^{1/5}\right)$ evaluations due to Aaronson [quant-ph/0111102] , and is tight up to a constant factor.

Our result also implies a quantum lower bound of $\Omega\left(n^{2/3}\right)$ queries to the inputs for the element distinctness problem, which is to determine whether or not the given n real numbers are distinct. The previous best lower bound is $\Omega\left(\sqrt{n}\right)$ queries in the black-box model; and $\Omega\left(\sqrt{n}\log n\right)$ comparisons in the comparisons-only model, due to Høyer, Neerbek, and Shi [ICALP'01, quant-ph/0102078].

Key words: Collision problem, element distinctness, lower bounds, quantum computation, computational complexity, polynomial method, quantum cryptography.

[1] This work was supported in part by the National Science Foundation under Grant. No. CCR-9820855, EIA-0086038, and CCR-0049092.

1 Introduction and summary of results

The exponential speed-up of Shor's quantum algorithm for integer factorization [21] over the best known classical algorithm has inspired scientists of many fields to explore the power of quantum computing. On the other hand, understanding the limitations of quantum computing is also of great importance. Identifying problems that are hard for quantum computers can not only deepen our knowledge on the power of quantum computing, but is also necessary for developing a new cryptography immune to quantum cryptanalysis.

Given a function f as an oracle, the collision problem is to find two distinct inputs i and j such that $f(i) = f(j)$, under the promise that such inputs exist. This paper concerns the r-to-one collision problem, in which the oracle is promised to be r-to-one, for some integer r fixed in advance. The case $r = 2$ is important because random two-to-one functions are considered good models of *collision intractable functions*, which is a fundamental cryptographic primitive. An exponential (in $\log n$) quantum lower bound would be evidence for the existence of collision intractable functions for quantum computers.

Other motivations of our study arise from the close connection of our problem to other widely-studied problems. An example is the *hidden subgroup problem*, in which the input is some r-to-one function with additional promises. The Abelian case of the hidden subgroup problem can be solved efficiently by a natural generalization of the well-known quantum algorithms of Simon [23] and Shor [22], while the non-Abelian case is one of the major challenges in the design of fast quantum algorithms (refer to Grigni, Schulman, Vazirani, and Vazirani [13] for a rencent development). A quantum lower bound for Collision would illuminate our understanding of the problem structures that allow or disallow a quantum speed-up.

It is not hard to see that $\Theta(\sqrt{n/r})$ evaluations are sufficient and necessary for classical algorithms to solve the r-to-one Collision. Interestingly, quantum computers can do much better: using Grover's quantum search algorithm [15] in a novel way, the quantum algorithm found by Brassard, Høyer, and Tapp [9] makes only $O\left((n/r)^{1/3}\right)$ evaluations. Despite much research effort, no lower bound better than constant had been found until very recently, when Aaronson proved the ground-breaking $\Omega\left((n/r)^{1/5}\right)$ lower bound [1]. In this paper, we improve the lower bound to the tight bound.

Theorem 1.1 (Lower bound for Collision). *Let $n > 0$ and $r \geq 2$ be integers with $r|n$, and let a function of domain size n be given as an oracle with the promise that it is either one-to-one or r-to-one. Then any error-bounded quantum algorithm for distinguishing these two cases must evaluate the function $\Omega\left((n/r)^{1/3}\right)$ times. Thus, finding a collision in an r-to-one function of domain size n requires $\Omega\left((n/r)^{1/3}\right)$ evaluations.*

Denote the set $\{1, 2, \cdots, n\}$ by $[n]$. It remains an open problem whether or not our lower bound still holds if the range of oracle is restricted to $[n]$. This is because Theorem 1.1 is proved by considering oracles with range $[3n/2]$. Nevertheless, for the small range case, we are able to improve Aaronson's $\Omega\left((n/r)^{1/5}\right)$ lower bound [1] to $\Omega\left((n/r)^{1/4}\right)$.

Theorem 1.2 (Lower bound for Collision with small range). *Let $n > 0$ and $r \geq 2$ be integers with $r|n$, and a function from $[n]$ to $[n]$ is given as an oracle with the promise that it is either one-to-one or r-to-one. Then any quantum algorithm for distinguishing*

these two cases must evaluate the function $\Omega\left((n/r)^{1/4}\right)$ times. Thus, finding a collision in an r-to-one function from $[n]$ to $[n]$ must evaluate the function $\Omega\left((n/r)^{1/4}\right)$ times.

Given n real numbers, are they all distinct? This is the classical problem of Element Distinctness, studied by many authors in the classical setting. A simple algorithm would be to sort the numbers using $\Theta(n \log n)$ comparisons, and then check the equality of neighboring numbers. This is essentially optimal classically, as suggested by the many $\Omega(n \log n)$ lower bounds in various classical models. In contrast, with another creative use of Grover's algorithm [15], the quantum algorithm found by Buhrman, Dürr, Heiligman, Høyer, Magniez, Santha, and de Wolf [10]makes only $O(n^{3/4} \log n)$ comparisons. Collision and Element Distinctness are closely related, as we can see from the following well-known reduction:

Reduction 1.3 (From Two-to-one Collision to Element Distinctness). Run the algorithm for Element Distinctness on the restriction of the oracle function on a random set of $\Theta(\sqrt{n})$ inputs. If the oracle is two-to-one, a collision will be found with high probability, by the Birthday Paradox.

Therefore, Theorem 1.1 implies,

Corollary 1.4 (Lower bound for Element Distinctness). *Any quantum algorithm that accesses the inputs through an oracle and solves the element distinctness problem of n real numbers must make $\Omega\left(n^{2/3}\right)$ oracle queries. If only comparisons are allowed, the same number of comparisons are required.*

The previous best known quantum lower bound is $\Omega(\sqrt{n})$ queries to the inputs, which can be obtained by a simple reduction from the search problem; and $\Omega(\sqrt{n} \log n)$ comparisons in the comparisons-only model, due to Høyer, Neerbek, and Shi [16]. The gap between our lower bound and the $O\left(n^{3/4} \log n\right)$ upper bound of Buhrman et al. [10] remains to be closed. The strongest classical lower bound is the $\Omega(n \log n)$ lower bound on the depth of randomized algebraic decision trees, due to Grigoriev, Karpinski, Meyer auf der Heide, and Smolensky [14]. For classical lower bounds in weaker models refer to the papers by Ben-Or [7], Steele and Yao [24], and, Dobkin and Lipton [12].

Remark 1.5. The worse-case and average-case complexities of the collision problems considered here are the same because of their symmetry. The reader may find it helpful to regard the problems as bipartite graph properties, and the inputs as bipartite graphs.

2 Proof outline and previous works

2.1 Proof outline

From now on, we shall refer to distinguishing an r-to-one function from a one-to-one function as the r-to-one problem, and denote it by $D_{r\to 1}$, or $D_{r\to 1}(n, N)$ when the domain and range sizes are n and N, respectively. For simplicity, we shall deal with $r = 2$ in this section.

Our proof for Theorem 1.1 takes two steps: first we reduce to $D_{2\to 1}$ a new problem *Half-two-to-one*, which is then shown to have an $\Omega\left(n^{1/3}\right)$ lower bound. Denote the set $\left\{\frac{n}{2}+1, \frac{n}{2}+2, \cdots, n\right\}$ by $[\frac{n}{2}+]$ (n is even).

Definition 2.1. Let $n > 0$ be an integer and $4|n$. In the **half-two-to-one problem**, or $D^{1/2}_{2\to 1}(n,n)$ for short, a function from $[n]$ to $[n]$ is given as an oracle with the promise that half of the inputs are two-to-one mapped to $[\frac{n}{2}+]$, and the other half are mapped to $[\frac{n}{2}]$, either one-to-one or two-to-one. The problem is to distinguish these two cases.

Lemma 2.2. *$D^{1/2}_{2\to 1}(n,n)$ can be reduced to $D_{2\to 1}(n, 3n/2)$ with a constant factor slow-down.*

Theorem 2.3. *Any quantum algorithm for $D^{1/2}_{2\to 1}(n,n)$ requires $\Omega(n^{1/3})$ evaluations.*

The reduction is done by exploring the symmetry of the problems, and by using the following important fact: on the $n/2$ inputs mapped to $[\frac{n}{2}+]$, f can be modified to be one-to-one mapped to $[3n/2]\backslash[n/2]$, without much slow-down.

We prove Theorem 2.3 by using the polynomial method of Beals, Buhrman, Cleve, Mosca, and de Wolf [6], and Aaronson [1] with new ideas. More specifically, let us fix a T-queries algorithm $\mathcal{A}$ for $D^{1/2}_{2\to 1}(n,n)$. First we *symmetrize* $\mathcal{A}$ to obtain $\bar{\mathcal{A}}$ so that running $\bar{\mathcal{A}}$ on any input f is equivalent to running $\mathcal{A}$ on a random input $\boldsymbol{f}$ "isomorphic" to f. Then we run $\bar{\mathcal{A}}$ on the oracle $f_{m,g}$, the function g-to-one mapped to $[n/2]$ on the first m inputs and two-to-one mapped to $[\frac{n}{2}+]$ on the remaining.

Following an important observation of Beals et al. [6] that relates the number of quantum queries to polynomial degrees, and from the nice symmetry of $\bar{\mathcal{A}}$, the acceptance probability $\bar{P}(f_{m,g})$ turns out to be a polynomial in m and g with degree $\leq 2T$. In addition, for all m and g such that $f_{m,g}$ is well-defined, $\bar{P}(f_{m,g}) \in [0,1]$; and, there is a gap between $\bar{P}(f_{\frac{n}{2},1})$ and $\bar{P}(f_{\frac{n}{2},2})$. These two nice properties enable one to apply a theorem by Paturi [19] to prove the desired lower bound for $\deg(\bar{P}(f_{m,g}))$. We point out that essentially Paturi's theorem follows from both Markov Inequality and Bernstein Inequality, two fundamental theorems in approximation theory that give good lower bounds for polynomial degrees.

For proving Theorem 1.2, we need the following additional idea. Given an algorithm for $D_{2\to 1}(n,n)$, we modify the algorithm so that it can be run on inputs that are only *partially* defined: Whenever the algorithm queries an undefined input, we force the algorithm to *abort* on the corresponding base vector. The rest of the proof is similar to that for Theorem 1.1.

Remark 2.4. Running the symmetrized algorithm on a fixed input is equivalent to running the algorithm on some random input, as treated by Aaronson [1]. However, we feel that our treatment explores the symmetry of the problem more explicitly and thus makes it less mysterious that the acceptance probability turns out to be a polynomial.

2.2 Relation with previous works

Aaronson [1] introduces the following original lower bound idea, which we shall refer to as the *derived polynomial method*: run the given T-queries algorithm on $\boldsymbol{f}_y$, a probability distribution determined by a parameter y. A new polynomial on y of $O(T)$ degree is derived from the average acceptance probability, and is then shown to have high degree by other methods. He is also the first to consider running the given algorithm on almost g-to-one functions for arbitrary g. We follow this approach in proving Theorem 2.3 and

improve his proof in the following ways: (1) The derived polynomial method seems to be more effective on Half-two-to-one than on Two-to-one itself. This is because the structure of Half-two-to-one yields a polynomial that has a gap around $m = n/2$, while the range of m is $[0, n]$. This feature, lacking in [1], is very important because it allows one to apply Bernstein Inequality, which in general gives a better degree lower bound than Markov Inequality if the function value has a sudden change close to the center of the domain. (2) The corresponding input distributions in our proofs are more natural and effective. As consequences, not only the ranges of the parameters are larger, but also the acceptance probabilities are exactly polynomials, instead of being close to a polynomial as in [1]. Thus better lower bounds can be obtained with simpler algebra.

It seems to us that our partial input idea was not used before. Another novel way of manipulating inputs for proving quantum lower bounds is used by Ambainis [3], where an adaptive adversary changes the input according to the performance of the algorithm.

Our problem can be formulated in the *black-box computation model*, a model widely studied in recent years due to both its simplicity and its power in modeling many natural problems. For other techniques for proving quantum lower bounds in this model, and quantum black-box computation in general, refer to the excellent survey of Ambainis [5].

We remark that previous approaches for proving degree lower bounds for (partial) Boolean functions can be interpreted in the light of the derived polynomial method. For example, the *symmetrization method*, introduced by Minsky and Papert [17] and used by Paturi [19] and Nisan and Szegedy [18], symmetrizes a Boolean function uniformly over all permutations of the Boolean variables. Another example, the *linear approximation* technique used by Shi [20], averages a Boolean function by tossing independent coins for each Boolean variable, and the mean value of each coin is a linear function of a single parameter.

The rest of this paper is organized as follows. In Section 3, we define the black-box model, introduce some notations, and state theorems from approximation theory which our proofs finally rely on the theorem of Paturi [19]. We then prove our lower bound for the general case of Collision in Section 4, which is followed by the proof for the special case of small range. Finally, we discuss some open problems.

3 Preparations for the proofs

Let $n \geq 0$ and $N \geq 0$ be integers and $\mathcal{F} := \mathcal{F}(n, N)$ be the set of all functions from $[n]$ to $[N]$. Let $f \in \mathcal{F}$ be given as an oracle. Following Beals et al. [6], we give the following definition of the black-box model, customized to our setting.

A quantum black-box algorithm works in a Hilbert space of dimension $n^2 L$, for some $L := L(n) < +\infty$. An orthonormal basis is chosen and denoted by

$$\{|i\rangle|j\rangle|l\rangle : i, j \in [n], l \in [L]\}.$$

For $j \in [N]$ and $j' \in [N]$, define $j + j' \mod N := i + j - \lfloor (i + j + 1)/N \rfloor \cdot N$. An oracle gate is the following unitary operator determined by f:

$$\mathsf{O}_f|i, j, l\rangle := |i,\ f(i) + j \mod N,\ l\rangle, \qquad \forall i \in [n], j \in [N], l \in [L].$$

A quantum black-box algorithm that makes T queries consists of $T + 1$ unitary operators, $\mathsf{U}_0, \mathsf{U}_1, \cdots, \mathsf{U}_T$, and a projection operator P, on the Hilbert space. It starts with a constant

vector denoted by $|0\rangle$, then applies the following sequence of operators:

$$\mathsf{U}_0 \to \mathsf{O}_f \to \mathsf{U}_1 \to \cdots \to \mathsf{U}_{T-1} \to \mathsf{O}_f \to \mathsf{U}_T \to \mathsf{P}.$$

The acceptance probability is

$$P(f) := \left\| \mathsf{P}\mathsf{U}_T\mathsf{O}_f\mathsf{U}_{T-1}\cdots\mathsf{O}_f\mathsf{U}_0|0\rangle \right\|^2.$$

We say that the algorithm computes a function $\phi : \mathcal{F} \supseteq \mathcal{F}' \to \{0,1\}$, where $\mathcal{F}' \subseteq \mathcal{F}$, with error probability bounded by ϵ if for every $f \in \mathcal{F}'$, $|P(f) - \phi(f)| \leq \epsilon$. The quantum complexity of ϕ is the minimal integer T such that there exists a quantum algorithm that computes ϕ with T queries and errs with a probability bounded by $1/3$.

As before, for all $i \in [n]$ and $j \in [N]$, the predicate $\delta_{i,j}(f) := 1$ if and only if $f(i) = j$. Observe that for all i, j, l, and f,

$$\mathsf{O}_f|i,j,l\rangle = \sum_{j'=1}^{N} \delta_{i,j'}(f)\,|i,\ j+j' \mod N,\ l\rangle.$$

Since all U_t and P are linear transformations, we have the following important observation by Beals et al. [6], in the form stated in Aaronson [1]:

Lemma 3.1. *The acceptance probability $P(f)$ can be expressed as a polynomial over the predicates $\delta_{i,j}$, $i \in [n], j \in [N]$, and $\deg(P) \leq 2T$.*

Let $\mathcal{F}^* := \mathcal{F}^*(n,N)$ denote the set of all *partial* functions from $[n]$ to $[N]$. Denote the domain and image of a function f^* by $\mathrm{dom}(f^*)$ and $\mathrm{img}(f^*)$, respectively. Any $f^* \in \mathcal{F}^*$ can be conveniently represented as a subset of $[n] \times [N]$, i.e., $f^* = \{(i, f^*(i)) : i \in \mathrm{dom}(f^*)\}$. For a finite set $K \subseteq \mathbb{Z}^+$, let $SG(K)$ denote the group of permutations on K. Any permutation in $SG(K)$ is understood as the identity mapping on any $k' \notin K$. For any integer $k > 0$, $SG(k)$ is a shorthand for $SG([k])$. For each $\sigma \in SG(n)$ and $\tau \in SG(N)$, define $\Gamma_\tau^\sigma : \mathcal{F}^* \to \mathcal{F}^*$ as

$$\Gamma_\tau^\sigma(f^*) := \{(\sigma(i), \tau(j)) : (i,j) \in f^*\}, \qquad \forall f^* \in \mathcal{F}^*.$$

For all $s \in \mathcal{F}^*$, the predicate $I_s : \mathcal{F}^* \to \{0,1\}$ is defined as follows:

$$I_s(f^*) := 1 \iff s \subseteq f^*, \qquad \forall f^* \in \mathcal{F}^*.$$

Fix a quantum black-box algorithm that queries T times. By Lemma 3.1, the acceptance probability can be written as

$$P(f) = \sum_{s \in \mathcal{F}^*, \mathsf{card}(s) \leq 2T} \beta_s\ I_s(f), \qquad \forall s,\ \beta_s \in \mathbb{R}. \tag{1}$$

Now proving a quantum lower bound is reduced to proving a lower bound on $\deg(P)$, for which we will resort to the following two fundamental theorems from approximation theory. For any function $q : \mathbb{R} \to \mathbb{R}$, and any set $D \subseteq \mathbb{R}$, let $\|q\|_D$ denote $\sup\{|q(\alpha)| : \alpha \in D\}$.

Theorem 3.2 (Markov Inequality). *For any polynomial $q(\alpha) \in \mathbb{R}[\alpha]$ with degree d and $\|q\|_{[-1,1]} = 1$,*

$$\|q'\|_{[-1,1]} \leq d^2.$$

Theorem 3.3 (Bernstein Inequality). *For any polynomial $q(\alpha) \in \mathbb{R}[\alpha]$ with degree d and $\|q\|_{[-1,1]} = 1$,*

$$|q'(\alpha)| \leq \frac{d}{\sqrt{1-\alpha^2}}, \qquad \forall \alpha \in (-1,1).$$

The proofs for the above theorems can be found in Chapter 4 of the book by Devore and Lorentz [11]. We will actually use the following result that follows from the above theorems. It is proven (with slight modification) by Paturi [19] in giving tight bounds for the lowest degree polynomial approximation to symmetric Boolean functions.

Theorem 3.4 (Paturi [19]). *Let $q(\alpha) \in \mathbb{R}[\alpha]$ be a polynomial of degree d, a and b be integers with $a < b$, and $\xi \in [a,b]$ be a real number. If (1) $|q(i)| \leq 1$ for all integers $i \in [a,b]$; and, (2) $|q(\lfloor \xi \rfloor) - q(\xi)| \geq c$ for some constant $c > 0$. Then,*

$$d = \Omega\left(\sqrt{(\xi - a + 1)(b - \xi + 1)}\right).$$

In particular,

$$d = \Omega\left(\sqrt{b-a}\right).$$

As a convention, all random variables are uniform over their domain.

4 Lower bound for the general collision problem

4.1 The reduction

Proof of Lemma 2.2. Let $\mathcal{A}$ be a quantum algorithm for $D_{2\to 1}(n, 3n/2)$. We shall derive an algorithm $\mathcal{B}$ for $D^{1/2}_{2\to 1}(n,n)$.

We call a function f half-two-to-one, if it is one-to-one on a half of its input, two-to-one on the other half, and, the two images are disjoint. Let $p_1 \geq 2/3$, $p_0 \leq 1/3$, and $p_{1/2}$ be the acceptance probabilities of $\mathcal{A}$ with the input being a random two-to-one, one-to-one, and half-two-to-one function from $[n]$ to $[3n/2]$, respectively. Let f be the oracle function for the $D^{1/2}_{2\to 1}(n,n)$ problem. Then f is either half-two-to-one or two-to-one, with some additional constrains on the range.

If $p_{1/2} < 1/2$, $\mathcal{B}$ will be the following: Choose random variables $\boldsymbol{\sigma} \in SG[n]$ and $\boldsymbol{\tau} \in SG[3n/2]$, then run $\mathcal{A}$ on $\boldsymbol{f} := \Gamma^{\boldsymbol{\sigma}}_{\boldsymbol{\tau}}(f)$. If f is two-to-one, the algorithm will accept with probability $p_1 \geq 2/3$; otherwise it will accept with probability $p_{1/2} < 1/2$.

Assume $p_{1/2} \geq 1/2$. Define $\bar{f} : [n] \to [3n/2]$ as:

$$\bar{f}(i) := \begin{cases} i + n/2 & \text{if } f(i) > n/2; \\ f(i) & \text{otherwise.} \end{cases} \tag{2}$$

Notice that the oracle $\mathsf{O}_{\bar{f}}$ can be simulated by two applications of O_f together with some local unitary operators. Now $\mathcal{B}$ will be: Choose random variables $\boldsymbol{\sigma} \in SG(n)$, and $\boldsymbol{\tau} \in SG(3n/2)$, then run $\mathcal{A}$ on $\boldsymbol{f} := \Gamma^{\boldsymbol{\sigma}}_{\boldsymbol{\tau}}(\bar{f})$.

Note that for each i with $f(i) \in [\frac{n}{2}+]$, $\bar{f}(i)$ is a distinct number in $[3n/2]\backslash[n/2]$. Therefore, if f is half-two-to-one, $\bar{f}$ is one-to-one, in which case $\boldsymbol{f}$ is a random one-to-one function; thus $\mathcal{B}$ will accept with probability $p_0 \leq 1/3$. On the other hand, if f is two-to-one, $\bar{f}$ is half-two-to-one, in which case $\boldsymbol{f}$ is a random half-two-to-one function; thus $\mathcal{B}$ will accept with probability $p_{1/2} \geq 1/2$. □

4.2 Lower bound for the half-two-to-one problem

Fix a quantum algorithm for $D^{1/2}_{2\to 1}(n,n)$, and let $P(f)$ be its acceptance probability. To prove an $\Omega\left(n^{1/3}\right)$ lower bound for $D^{1/2}_{2\to 1}(n,n)$, we need only to prove the lower bound for $\deg(P)$, by Lemma 3.1. Define the symmetrization of P as

$$\bar{P}(f) := \mathrm{E}_{\boldsymbol{\sigma}\in SG(n),\boldsymbol{\tau}\in SG([n/2]),\boldsymbol{\tau'}\in SG([\frac{n}{2}+])}\left[\ P\left(\Gamma^{\boldsymbol{\sigma}}_{\boldsymbol{\tau}\circ\boldsymbol{\tau'}}(f)\right)\ \right]. \tag{3}$$

Definition 4.1. We call a pair of integers (m,g) **valid**, if $0 \leq m \leq n$, $1 \leq g \leq n$, $2|m$, $g|m$, and if $g = 1$, $m \leq n/2$.

Given a valid (m,g), define $f_{m,g} : [n] \to [n]$ as follows:

$$f_{m,g}(i) = \begin{cases} \lceil i/g \rceil & i \in [m], \\ \lceil (i-m)/2 \rceil + n/2 & \text{otherwise.} \end{cases} \tag{4}$$

Lemma 4.2. *The function $\bar{P}(f_{m,g})$ is a polynomial in m and g of degree $\leq 2T$.*

Proof. By Lemma 3.1, it suffices to show that for each monomial I_s, $\mathsf{card}(s) \leq 2T$, the symmetrization $\bar{I}_s$ is such a polynomial, where

$$\bar{I}_s(f_{m,g}) := \mathrm{E}_{\boldsymbol{\sigma},\boldsymbol{\tau},\boldsymbol{\tau'}}\left[\ I_s\left(\Gamma^{\boldsymbol{\sigma}}_{\boldsymbol{\tau}\circ\boldsymbol{\tau'}}(f_{m,g})\right)\ \right].$$

Let $w := \mathsf{card}(\mathrm{img}(s) \cap [n/2])$. Fix a sequence of elements in $\mathrm{img}(s) \cap [n/2]$, and let $u_1, u_2, \cdots, u_w$ be the corresponding sequence of sizes of preimages for the elements. Put $u := \sum_{j=1}^{w} u_j$. Replacing $[n]$ by $[\frac{n}{2}+]$, we define w', u'_j, $1 \leq j \leq w'$, and u', similarly. For integers a, b, $P^b_a := a(a-1)\cdots(a-b+1)$. Put

$$\lambda := \frac{(n/2-w)!(n/2-w')!(n-u-u')!}{n!(n/2)!(n/2)!}.$$

By simple calculations,

$$\begin{aligned} \bar{I}_s(f_{m,g}) &= \lambda \cdot P^w_{m/g} \cdot \Pi^w_{j=1} P^{u_i}_g \cdot P^{w'}_{\frac{n-m}{2}} \cdot \Pi^{w'}_{j=1} P^{u'_i}_2 && (5) \\ &= \lambda \cdot \Pi^{w-1}_{j=0}(m - g\cdot j) \cdot \Pi^w_{j=1} P^{u_i-1}_{g-1} \cdot \Pi^{w'-1}_{j=0}(n-m-2j) \cdot \Pi^{w'}_{j=1} P^{u'_i-1}_1, && (6) \end{aligned}$$

which is a polynomial in m and g of degree

$$w + (u-w) + w' + (u'-w') = u + u' = \mathsf{card}(s) \leq 2T.$$

□

Proof of Theorem 2.3. Since $\deg(\bar{P}(f_{m,g})) \leq 2T$ by the above lemma, it suffices to prove $\deg(\bar{P}(f_{m,g})) = \Omega\left(n^{1/3}\right)$.

Since $\bar{P}(f_{m,g})$ is defined to be the acceptance probability for the oracle $f_{m,g}$,

$$0 \leq \bar{P}(f_{m,g}) \leq 1, \quad \text{for all valid } (m,g), \tag{7}$$

$$0 \leq \bar{P}(f_{n/2,1}) \leq 1/3, \qquad \text{and,} \qquad 2/3 \leq \bar{P}(f_{n/2,2}) \leq 1. \tag{8}$$

Put $G := \lfloor n^{2/3} \rfloor$, and $Q_1(\alpha) := \bar{P}(f_{n/2,\alpha})$. Clearly, $\deg(Q_1) \leq \deg(\bar{P}(f_{m,g}))$. By Equations in 8,

$$|Q_1(1) - Q_1(2)| = \left|\bar{P}(f_{n/2,1}) - \bar{P}(f_{n/2,2})\right| \geq 1/3.$$

If $|Q_1(k)| \leq 2$ for all $k \in [G]$, by Theorem 3.4, $\deg(Q_1) = \Omega\left(\sqrt{G}\right)$, which implies $\deg(\bar{P}(f_{m,g})) = \Omega\left(n^{1/3}\right)$. Otherwise, let $g_0 \in [G]$ be such that $|Q_1(g_0)| > 2$.

Put $G_0 := \lfloor \frac{n}{2g_0} \rfloor$, and,

$$Q_2(\alpha) := \bar{P}(f_{2g_0\alpha,\, g_0}), \qquad \alpha \in [0, G_0].$$

Then $G_0 = \Omega\left(n^{1/3}\right)$, and $\deg(Q_2) \leq \deg(Q)$. Since $g_0 \geq 2$, $(2ig_0, g_0)$ is valid for each $i \in [G_0^*]$, which implies $0 \leq Q_2(i) \leq 1$, by Eqn. 7. Since

$$\left|Q_2\left(\frac{n}{4g_0}\right)\right| = \left|\bar{P}(f_{n/2,\, g_0})\right| = |Q_1(g_0)| > 2,$$

and $0 \leq Q_2\left(\left\lfloor \frac{n}{4g_0} \right\rfloor\right) \leq 1$, we have,

$$\left|Q_2\left(\left\lfloor \frac{n}{4g_0} \right\rfloor\right) - Q_2\left(\frac{n}{4g_0}\right)\right| \geq 1.$$

Applying Theorem 3.4, we have

$$\deg(Q_2) = \Omega\left(\sqrt{\left(\frac{n}{4g_0} + 1\right)\left(G_0 - \frac{n}{4g_0} + 1\right)}\right),$$

which implies $\deg(\bar{P}(f_{m,g})) = \Omega\left(n^{1/3}\right)$. □

4.3 Generalizing to arbitrary $r \geq 2$

Proof of Theorem 1.1. Combining Lemma 2.2 and Theorem 2.3 we obtain Theorem 1.1 for the case $r = 2$. To generalize to arbitrary $r \geq 2$, we need only to replace Half-two-to-one by *Half-r-to-one*, denoted by $D^{1/2}_{r\to 1}(n, \frac{n}{2} + \frac{n}{r})$, where the oracle is r-to-one mapped to $[n/2+1, n/2+2, \cdots, n/2+n/r]$ on $n/2$ inputs and the other $n/2$ inputs are mapped to $[n/2]$ either r-to-one or one-to-one. In Definition 4.1, the condition $2|m$ for (m,g) being valid is replaced by $r|m$.

In analogy to Lemma 2.2, $D^{1/2}_{r\to 1}(n, \frac{n}{2} + \frac{n}{r})$ can be reduced to $D_{r\to 1}(n, 3n/2)$. To prove the $\Omega\left((n/r)^{1/3}\right)$ lower bound for the former, we need only to modify the proof for the latter by choosing appropriate parameters. That is, we set $G := \left(\lfloor (n/r)^{2/3} \rfloor\right) \cdot r$. We leave the remaining work for interested readers. □

5 Lower bound for Collision with small range

Let n and r be integers, and $r|n$. Fix a T-queries quantum black-box algorithm for $D_{r\to 1}(n,n)$. Let $P(f)$ be its acceptance probability. Instead of making a reduction, we need the following lemma.

Lemma 5.1. *For any partial assignment s,*

$$0 \le P(s) \le 1.$$

Proof. Let U_t, $0 \le t \le T$, and P be the unitary operators and the final projection operator of the algorithm. Let P_s be the operator that projects a state to the subspace spanned by

$$\{|i,j,l\rangle : i \in \mathrm{dom}(s), j \in [n], l \in [L]\}.$$

Then it can be easily proved by induction that

$$P(s) = \|\mathsf{P}\mathsf{P}_s\mathsf{U}_T\mathsf{O}_x\mathsf{P}_s \cdots \mathsf{P}_s\mathsf{U}_1\mathsf{O}_x\mathsf{P}_s\mathsf{U}_0|0\rangle\|^2.$$

The lemma follows. □

The symmetrization of P is defined as

$$\bar{P}(f) := \mathrm{E}_{\boldsymbol{\sigma},\boldsymbol{\tau}\in SG(n)}\left[P(\Gamma^{\boldsymbol{\sigma}}_{\boldsymbol{\tau}}(f))\right].$$

Now we call a pair of integers (m,g) *valid* if $m \in [n^*]$, $g \in [n]$, and $g|m$. Given a valid (m,g), define the partial function $f_{m,g}$ as follows:

$$f_{m,g} := \{(i, \lceil i/g \rceil) : i \in [m]\}.$$

By Lemma 5.1 and the definition of $\bar{P}(f_{m,g})$,

$$0 \le \bar{P}(f_{m,g}) \le 1, \qquad \text{for all valid } (m,g). \tag{9}$$

By the correctness of the algorithm,

$$2/3 \le \bar{P}(f_{n,1}) \le 1, \qquad \text{and,} \qquad 0 \le \bar{P}(f_{n,r}) \le 1/3. \tag{10}$$

Lemma 5.2. *The function $\bar{P}(f_{m,g})$ can be expressed as a polynomial in m and g of degree $\le 2T$.*

We omit the proof since it is in analogy to the proof for Lemma 4.2.

Proof of Theorem 1.2. By Lemma 5.2, it suffices to prove $\deg(\bar{P}(f_{m,g})) = \Omega((n/r)^{1/4})$. The proof is similar to that for Theorem 2.3, and is much simpler. We leave the details to the reader. □

6 Open problems

Besides the two mentioned open problems, Collision with small range and Element Distinctness, we raise two more.

Definition 6.1. Two sets $f = \{f(1), f(2), \cdots, f(n)\}$ and $g = \{g(1), g(2), \cdots, g(n)\}$ are given as oracles with the promise that either $f = g$ or $f \cap g = \varnothing$. The **set equality problem** is to distinguish these two cases.

This is a special case of the two-to-one problem, and it closely models the Graph Isomorphism problem. We are not able to prove any $\omega(1)$ lower bound, while we conjecture that it is as hard as the general Collision. A problem harder than the above is:

Definition 6.2. Given n distinct numbers $x_1, x_2, \ldots, x_n$, the **index erasure problem** is to generate a vector close to $|\phi_x\rangle = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} |x_i\rangle$.

This problem is equivalent to the following **quantum-parallel search problem**: Given an oracle described above, and the state $|\phi_x\rangle$, generate a vector close to $\frac{1}{\sqrt{n}} \sum_{i=1}^{n} |x_i\rangle|i\rangle$. One can show that $O(\sqrt{n})$ queries are sufficient for both problems by using Grover's quantum search algorithm [15]. We conjecture that this is tight, though we are not able to prove any $\omega(1)$ lower bound.

7 Acknowledgments

I am most grateful to Ronald de Wolf and Scott Aaronson for stimulating discussions and valuable comments, and to Scott Aaronson for sending me his results on Collision. Other thanks go to Lawrence Ip, Ashwin Nayak, and Leonard Schulman for helpful discussions and comments. I am also indebted to Andy Yao, Umesh Vazirani, and Ashwin Nayak for introducing these problems to me and for the numerous discussions at the early stage of this work.

References

[1] S. Aaronson. Quantum lower bound for the collision problem. LANL e-print, quant-ph/0111102, 2001.

[2] *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, Philadelphia, Pennsylvania, May 1996.

[3] A. Ambainis. A better lower bound for quantum algorithms searching an ordered list. In *40th Symposium on Foundations of Computer Science (FOCS)*, pages 352–357, New York, NY, USA, Oct. 1999. IEEE Computer Society.

[4] A. Ambainis. Quantum lower bounds by quantum arguments. In *Proceedings of the Thirty-second Annual ACM Symposium on the Theory of Computing*, pages 636–643, Portland, Oregon, May 2000.

[5] A. Ambainis. Quantum query algorithms and lower bounds. In *Proceedings of Foundations of the Formal Sciences III*, September 2001.

[6] R. Beals, H. Buhrman, R. Cleve, M. Mosca, and R. de Wolf. Quantum lower bounds by polynomials. In *39th Annual Symposium on Foundations of Computer Science*, pages 352–361, Los Alamitos, CA, Nov. 1998. IEEE.

[7] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proceedings of the Fifteenth Annual ACM Symposium on the Theory of Computing*, pages 80–86, Boston, Massachusetts, Apr. 1983.

[8] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, Oct. 1997.

[9] G. Brassard, P. Høyer, and A. Tapp. Quantum algorithm for the collision problem. *ACM SIGACT News*, 28:14–19, 1997. Preprint: quant-ph/9705002.

[10] H. Buhrman, C. Dürr, M. Heiligman, P. Høyer, F. Magniez, M. Santha, and R. de Wolf. Quantum algorithms for element distinctness. In *Proceedings of Sixteenth IEEE conference on Computational Complexity*, 2001.

[11] R. A. DeVore and G. G. Lorentz. *Constructive approximation.* Springer-Verlag, Berlin, 1993.

[12] D. Dobkin and R. Lipton. A lower bound of $\frac{1}{2}n^2$ on linear search programs for the knapsack problem. *J. Comput. Syst. Sci.*, 16:413–417, 1978.

[13] M. Grigni, L. Schulman, M. Vazirani, and U. Vazirani. Quantum mechanical algorithms for the nonabelian hidden subgroup problem. In *Proceedings of Thirty-Third Annual ACM Symposium on Theory of Computing*, 2001.

[14] D. Grigoriev, M. Karpinski, F. Meyer auf der Heide, and R. Smolensky. A lower bound for randomized algebraic decision trees. In ACM [2], pages 612–619.

[15] L. K. Grover. A fast quantum mechanical algorithm for database search. In ACM [2], pages 212–219.

[16] P. Høyer, J. Neerbek, and Y. Shi. Quantum complexities of ordered searching, sorting, and element distinctness. In *Proceedings of 28th International Colloquium on Automata, Languages, and Programming*, 2001.

[17] M. Minsky and S. Papert. *Perceptrons.* MIT Press, Cambridge, MA, 1969.

[18] N. Nisan and M. Szegedy. On the degree of Boolean functions as real polynomials. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on the Theory of Computing*, pages 462–467, Victoria, British Columbia, Canada, 4–6 May 1992.

[19] R. Paturi. On the degree of polynomials that approximate symmetric Boolean functions (preliminary version). In *Proceedings of the Twenty-Fourth Annual ACM Symposium on the Theory of Computing*, pages 468–474, Victoria, British Columbia, Canada, May 1992.

[20] Y. Shi. Approximating linear restrictions of Boolean functions. Manuscript, 2001.

[21] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings: 35th Annual Symposium on Foundations of Computer Science, November 20–22, 1994, Santa Fe, New Mexico*, pages 124–134. IEEE Computer Society Press, 1994.

[22] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, Oct. 1997.

[23] D. R. Simon. On the power of quantum computation. *SIAM Journal on Computing*, 26(5):1474–1483, Oct. 1997.

[24] J. M. Steele and A. C.-C. Yao. Lower bounds for algebraic decision trees. *J. Algorithms*, 3(1):1–8, 1982.